

Scholar

Results 1 - 10 of about 679 for **pagesize table resize OR reorganize OR rebuild**. (0.24 seconds)

Did you mean: **page size table resize OR reorganize OR rebuild**

Online index rebuild - group of 3 »

All articles [Recent articles](#)

N Ponnekanti, H Kodavalla - ACM SIGMOD Record, 2000 - portal.acm.org

... a shared **table lock** on the **table** thereby making the **table** inaccessible to OLTP transactions, which may not be acceptable. Online index **rebuild** restores the ...

Cited by 3 - [Related Articles](#) - [Web Search](#) - [BL Direct](#)

DB2 UDB V7.1 Performance Tuning Guide

T Shirai, L Dilworth, R Reutlinger, S Kumar, D ... - IBM International Technical Support Organization, December, 2000 - polaris.com.tr

... ix **Tables**. Page 13. © Copyright IBM Corp. 2000 xi **Tables** ...

Cited by 5 - [Related Articles](#) - [View as HTML](#) - [Web Search](#)

Reorganizing News Web Pages for Mobile Users - group of 2 »

W Kim, E Hwang, W Kim - LECTURE NOTES IN COMPUTER SCIENCE, 2003 - Springer

... S as shown in Figure 3. Then our proposed system would **reorganize** the web ... Reduced size of **page size**. **Table 3** compares the file size and display size before and ...

[Related Articles](#) - [Web Search](#) - [BL Direct](#)

Word 2003—Tutorial II - group of 2 »

P Formatting, D Canvas, P Formatting, K Shortcuts - wwwlb.aub.edu.lb

... Rows and Columns • Move and **Resize Tables** • Moving the ... Part of a **Tables/** a Whole

Table ... Footnotes and Endnotes • Bookmarks • **Page Size** and Orientation ...

[View as HTML](#) - [Web Search](#)

A note on storage fragmentation and program segmentation - group of 2 »

B Randell - Communications of the ACM, 1969 - portal.acm.org

... or of expending system resources in order to **reorganize** the contents ... SA2 HAS'b'

BITS,WHERE 2 b= P(PAGE SIZE) SPA IS IN UNITS ... PAGE TABLE FOR SEGMENT SN I A SPA ...

Cited by 35 - [Related Articles](#) - [Web Search](#)

Email and postal

J Hong, G Chen - cis.ohio-state.edu

... the array when the block size is set to **page size**. ... of JIAJIA as such: Cachepages = 2048, **Pagesize**=8192(8K ... only at the cold startup period (**Table 4**), and few ...

[Related Articles](#) - [View as HTML](#) - [Web Search](#)

A heuristic file reorganization algorithm based on record clustering - group of 3 »

P Scheuermann, YC Park, E Omiecinski - BIT Numerical Mathematics, 1989 - Springer

... Let the buffer size be 2 pages and the **page size** be 4 records. ... If $N_p < \frac{\text{pagesize}}{2}$, bring ...

... by setting their NPGs to 0 and delete k from RS and the **COST table**. ...

Cited by 3 - [Related Articles](#) - [Web Search](#) - [Library Search](#)

Effective "static-graph" reorganization to improve locality in garbage-collected systems - group of 2 »

PR Wilson, MS Lam, TG Moher - Proceedings of the ACM SIGPLAN 1991 conference on ..., 1991 - portal.acm.org

... they hold are often much smaller than the **page size**. Using hash **tables** to implement

large namespaces — as is ... applying traversal algorithms to **reorganize** data ...

Cited by 63 - [Related Articles](#) - [Web Search](#)

[doc] **INSTRUCTIONS FOR AUTHORS** - group of 2 »

HM Hien, V Raghavan - gisws.media.osaka-cu.ac.jp
... You might also be able to **resize** a figure. ... preparing illustrations, please bear in mind the **page size**, the space ... in their intended location in the text (**Table 1** ...
Related Articles - View as HTML - Web Search


Making Statistical Tables in Windows: A Comparison of Data Muncher and SAS (PROC TABULATE and PROC ... - group of 3 »
KL Monti - The American Statistician, 1997 - questia.com
... serverDoc.getElementByld(srcName).innerHTML; **resize()**; } } function showHideMessage ...
One can change the **page size** prior to generating the **table** with its ...
Related Articles - Web Search - BL_Direct

Did you mean to search for: **page size** table resize OR reorganize OR rebuild

Goooooooooooooogle ►
Result Page: 1 2 3 4 5 6 7 8 9 10 **Next**

[Google Home](#) - [About Google](#) - [About Google Scholar](#)

THE ACM DIGITAL LIBRARY

 [Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used increase and page size and reorganize

Found 4,094 of 185,030

Sort results by

☒ Save results to a Binder

Try an [Advanced Search](#)

Display results

☐ Search Tips

Try this search in [The ACM Guide](#)
☐ Open results in a new window

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale ☐ ☐ ☐ ☐ ☐

1 [Effective clustering of complex objects in object-oriented databases](#)



Jia-Bing R. Cheng, A. R. Hurson

April 1991 **ACM SIGMOD Record , Proceedings of the 1991 ACM SIGMOD international conference on Management of data SIGMOD '91**, Volume 20 Issue 2

Publisher: ACM Press

Full text available:  pdf(1.08 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

2 [Evaluation of concurrent physical database reorganization through simulation modeling](#)



Lars Söderlund

September 1981 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1981 ACM SIGMETRICS conference on Measurement and modeling of computer systems SIGMETRICS '81**, Volume 10 Issue 3

Publisher: ACM Press

Full text available:  pdf(880.64 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The performance of a database system commonly deteriorates due to degradation of the database's physical data structure. The structure degradation is a consequence of the normal operations of a general database management system. When system performance has degraded below acceptable limits the database must be reorganized. In conventional, periodic reorganization the database, or part of it, is taken off line while the data structure is being reorganized. This paper pr ...


3 [Tradeoffs in supporting two page sizes](#)



Madhusudhan Talluri, Shing Kong, Mark D. Hill, David A. Patterson

April 1992 **ACM SIGARCH Computer Architecture News , Proceedings of the 19th annual international symposium on Computer architecture ISCA '92**, Volume 20 Issue 2

Publisher: ACM Press

Full text available:  pdf(1.18 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

As computer system main memories get larger and processor cycles-per-instruction (CPIs) get smaller, the time spent in handling translation lookaside buffer (TLB) misses could become a performance bottleneck. We explore relieving this bottleneck by (a) increasing the page size and (b) supporting two page sizes. We discuss how to build a TLB to support two page sizes and examine both alternatives experimentally with a dozen unprogrammed, user-mode traces for the SPARC architectur ...

4 [Effective "static-graph" reorganization to improve locality in garbage-collected](#)



systems

Paul R. Wilson, Michael S. Lam, Thomas G. Moher

May 1991 **ACM SIGPLAN Notices , Proceedings of the ACM SIGPLAN 1991 conference on Programming language design and implementation PLDI '91**, Volume 26
Issue 6

Publisher: ACM Press

Full text available: pdf(1.31 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)



5 Application restructuring and performance portability on shared virtual memory and hardware-coherent multiprocessors



Dongming Jiang, Hongzhang Shan, Jaswinder Pal Singh

June 1997 **ACM SIGPLAN Notices , Proceedings of the sixth ACM SIGPLAN symposium on Principles and practice of parallel programming PPOPP '97**, Volume 32
Issue 7

Publisher: ACM Press

Full text available: pdf(1.59 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)



The performance portability of parallel programs across a wide range of emerging coherent shared address space systems is not well understood. Programs that run well on efficient, hardware cache-coherent systems often do not perform well on less optimal or more commodity-based communication architectures. This paper studies this issue of performance portability, with the commodity communication architecture of interest being page-grained shared virtual memory. We begin with applications that per ...

6 Locality preserving dictionaries: theory & application to clustering in databases



Vijayshankar Raman

May 1999 **Proceedings of the eighteenth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems**

Publisher: ACM Press

Full text available: pdf(1.09 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)



7 On the cost of monitoring and reorganization of object bases for clustering



Carsten A. Gerlhof, Alfons Kemper, Guido Moerkotte

September 1996 **ACM SIGMOD Record**, Volume 25 Issue 3

Publisher: ACM Press

Full text available: pdf(606.93 KB) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)



Clustering is one of the most effective means to enhance the performance of object base applications. Consequently, many proposals exist for algorithms computing good object placements depending on the application profile. However, in an effective object base reorganization tool the clustering algorithm is only one constituent. In this paper, we report on our object base reorganization tool that covers all stages of reorganizing the objects: the application profile is determined by a monito ...

8 Storage reorganization techniques for matrix computation in a paging environment



Patrick C. Fischer, Robert L. Probert

July 1979 **Communications of the ACM**, Volume 22 Issue 7

Publisher: ACM Press

Full text available: pdf(900.05 KB) Additional Information: [full citation](#), [references](#), [citations](#)



Keywords: data reorganization, matrix multiplication, pagination, paging, transpose, virtual memory

9 Query evaluation techniques for large databases



Goetz Graefe

June 1993 **ACM Computing Surveys (CSUR)**, Volume 25 Issue 2

Publisher: ACM Press

Full text available: pdf(9.37 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate simple records, query-processi ...

Keywords: complex query evaluation plans, dynamic query evaluation plans, extensible database systems, iterators, object-oriented database systems, operator model of parallelization, parallel algorithms, relational database systems, set-matching algorithms, sort-hash duality

10 Database Reorganization—Principles and Practice



Gary H. Sockut, Robert P. Goldberg

December 1979 **ACM Computing Surveys (CSUR)**, Volume 11 Issue 4

Publisher: ACM Press

Full text available: pdf(1.89 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

11 Code management: Dynamic code management: improving whole program code locality in managed runtimes



Xianglong Huang, Brian T Lewis, Kathryn S McKinley

June 2006 **Proceedings of the 2nd international conference on Virtual execution environments VEE '06**

Publisher: ACM Press

Full text available: pdf(153.03 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Poor code locality degrades application performance by increasing memory stalls due to instruction cache and TLB misses. This problem is particularly an issue for large server applications written in languages such as Java and C# that provide just-in-time (JIT) compilation, dynamic class loading, and dynamic recompilation. However, managed runtimes also offer an opportunity to dynamically profile applications and adapt them to improve their performance. This paper describes a Dynamic Code Manage ...

Keywords: code generation, code layout, dynamic optimization, locality, performance monitoring, virtual machines

12 Linear hashing with separators—a dynamic hashing scheme achieving one-access



Per-Ake Larson

September 1988 **ACM Transactions on Database Systems (TODS)**, Volume 13 Issue 3

Publisher: ACM Press

Full text available: pdf(1.62 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

A new dynamic hashing scheme is presented. Its most outstanding feature is that any record can be retrieved in exactly one disk access. This is achieved by using a small amount of supplemental internal storage that stores enough information to uniquely determine the current location of any record. The amount of internal storage required is small: typically one byte for each page of the file. The necessary address computation,


insertion, and expansion algorithms are presented and the perform ...

13 Data page layouts for relational databases on deep memory hierarchies

Anastassia Ailamaki, David J. DeWitt, Mark D. Hill

November 2002 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 11 Issue 3

Publisher: Springer-Verlag New York, Inc.

Full text available:  pdf(593.86 KB) Additional Information: [full citation](#), [abstract](#), [index terms](#)

Relational database systems have traditionally optimized for I/O performance and organized records sequentially on disk pages using the N-ary Storage Model (NSM) (a.k.a., slotted pages). Recent research, however, indicates that cache utilization and performance is becoming increasingly important on modern platforms. In this paper, we first demonstrate that in-page data placement is the key to high cache performance and that NSM exhibits low cache utilization on modern platforms. Next, we ...

Keywords: Cache-conscious database systems, Disk page layout, Relational data placement

14 Empirical working set behavior



Juan Rodriguez-Rosell

September 1973 **Communications of the ACM**, Volume 16 Issue 9

Publisher: ACM Press

Full text available:  pdf(457.79 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

The working set model for program behavior has been proposed in recent years as a basis for the design of scheduling and paging algorithms. Although the words "working set" are now commonly encountered in the literature dealing with resource allocation, there is a dearth of published data on program working set behavior. It is the purpose of this paper to present empirical data from actual program measurements, in the hope that workers in the field might find experimental evidence ...

Keywords: paging, program behavior, software measurement, virtual memory, working set


15 An improved network clustering method for I/O-efficient query processing



Sung-Ho Woo, Sung-Bong Yang

November 2000 **Proceedings of the 8th ACM international symposium on Advances in geographic information systems**

Publisher: ACM Press

Full text available:  pdf(651.16 KB) Additional Information: [full citation](#), [abstract](#), [index terms](#)

Efficient network query processing is extremely important in Geographical Information Systems (GIS) and Intelligent Transportation Systems (ITS) which include various applications of transportation, utility and communication networks, etc. In order to reduce the I/O cost in network query processing a given network should be stored with high disk-space utilization and a low edge-cut ratio. To do so the nodes in the network should be clustered in such a way that each cluster fits in a disk page ...

16 Eliminating the address translation bottleneck for physical address cache



Tzi-cker Chiueh, Randy H. Katz

September 1992 **ACM SIGPLAN Notices , Proceedings of the fifth international conference on Architectural support for programming languages and operating systems ASPLOS-V**, Volume 27 Issue 9

Publisher: ACM Press

Full text available:  pdf(1.28 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

17 Research sessions: implementation techniques: Fractal prefetching B⁺-Trees:
optimizing both cache and disk performance

Shimin Chen, Phillip B. Gibbons, Todd C. Mowry, Gary Valentin

June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on
Management of data SIGMOD '02**

Publisher: ACM Press

Full text available:  pdf(1.49 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)


B⁺-Trees have been traditionally optimized for I/O performance with disk pages as tree nodes. Recently, researchers have proposed new types of B⁺-Trees optimized for CPU cache performance in main memory environments, where the tree node sizes are one or a few cache lines. Unfortunately, due primarily to this large discrepancy in optimal node sizes, existing disk-optimized B⁺-Trees suffer from poor cache performance while cache-optimized B⁺-Trees exhibit ...

18 Modeling on-line rebalancing with priorities and executing on parallel database systems

Daniel C. Zilio

November 1996 **Proceedings of the 1996 conference of the Centre for Advanced
Studies on Collaborative research**

Publisher: IBM Press

Full text available:  pdf(213.63 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Because changes to the database (DB) and workload occur during a DB system's lifetime, the physical DB design must evolve to sustain good performance. These changes are carried out by on-line reorganizations which access the DB and execute concurrently with the DB workload. Different performance intrusions are placed on the workload when a reorganization is assigned different priorities compared to the workload processes. Our work studies the effects of the reorganization priority-level on performance ...

19 Physical storage structures: The K-D-B-tree: a search structure for large
multidimensional dynamic indexes

John T. Robinson

April 1981 **Proceedings of the 1981 ACM SIGMOD international conference on
Management of data SIGMOD '81**

Publisher: ACM Press

Full text available:  pdf(723.91 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)


The problem of retrieving multikey records via range queries from a large, dynamic index is considered. By *large* it is meant that most of the index must be stored on secondary memory. By *dynamic* it is meant that insertions and deletions are intermixed with queries, so that the index cannot be built beforehand. A new data structure, the *K-D-B-tree*, is presented as a solution to this problem. K-D-B-trees combine properties of K-D-trees and B-trees. It is expected that the multikey ...

20 Dynamic hashing schemes

R. J. Enbody, H. C. Du

July 1988 **ACM Computing Surveys (CSUR)**, Volume 20 Issue 2

Publisher: ACM Press

Full text available:  pdf(2.52 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

A new type of dynamic file access called dynamic hashing has recently emerged. It promises the flexibility of handling dynamic files while preserving the fast access times expected from hashing. Such a fast, dynamic file access scheme is needed to support modern database systems. This paper surveys dynamic file access schemes and examines their critical design issues.



Keywords: dynamic hashing

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)

Terms used **Adaptive record clustering**

Found 12,928 of 184,245

Sort results by

relevance

Display results

expanded form

 [Save results to a Binder](#)
 [Search Tips](#)
☐ Open results in a new window

Try an [Advanced Search](#)

Try this search in [The ACM Guide](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale ☐ ☐ ☐ ☐ ☐


1 [Adaptive record clustering](#)



C. T. Yu, Cheing-mei Suen, K. Lam, M. K. Siu

June 1985 **ACM Transactions on Database Systems (TODS)**, Volume 10 Issue 2

Publisher: ACM Press

Full text available:  pdf(1.58 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

An algorithm for record clustering is presented. It is capable of detecting sudden changes in users' access patterns and then suggesting an appropriate assignment of records to blocks. It is conceptually simple, highly intuitive, does not need to classify queries into types, and avoids collecting individual query statistics. Experimental results indicate that it converges rapidly; its performance is about 50 percent better than that of the total sort method, and about 100 percent better than ...


2 [A parallel algorithm for record clustering](#)



Edward Omiecinski, Peter Scheuermann

December 1990 **ACM Transactions on Database Systems (TODS)**, Volume 15 Issue 4

Publisher: ACM Press

Full text available:  pdf(1.82 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

We present an efficient heuristic algorithm for record clustering that can run on a SIMD machine. We introduce the P-tree, and its associated numbering scheme, which in the split phase allows each processor independently to compute the unique cluster number of a record satisfying an arbitrary query. We show that by restricting ourselves in the merge phase to combining only sibling clusters, we obtain a parallel algorithm whose speedup ratio is optimal in the number of processors used. Final ...


3 [Adaptive document clustering](#)



C. T. Yu, Y. T. Wang, C. H. Chen

June 1985 **Proceedings of the 8th annual international ACM SIGIR conference on Research and development in information retrieval**

Publisher: ACM Press

Full text available:  pdf(465.87 KB)

Additional Information: [full citation](#), [references](#), [citations](#)

4 [Adaptive information system design: one query at a time](#)



C. T. Yu, C. H. Chen

May 1985 **ACM SIGMOD Record, Proceedings of the 1985 ACM SIGMOD international conference on Management of data SIGMOD '85**, Volume 14 Issue 4

Publisher: ACM Press

Full text available:  [pdf\(1.03 MB\)](#)

Additional Information: [full citation](#), [references](#), [index terms](#)


5 Research articles and surveys: Research issues in automatic database clustering



Sylvain Guinepain, Le Gruenwald

March 2005 **ACM SIGMOD Record**, Volume 34 Issue 1

Publisher: ACM Press

Full text available:  [pdf\(1.42 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)


While a lot of work has been published on clustering of data on storage medium, little has been done about automating this process. This is an important area because with data proliferation, human attention has become a precious and expensive resource. Our goal is to develop an automatic and dynamic database clustering technique that will dynamically re-cluster a database with little intervention of a database administrator (DBA) and maintain an acceptable query response time at all times. In th ...

6 A global approach to record clustering and file reorganization

Edward Omiecinski, Peter Scheuermann

July 1984 **Proceedings of the 7th annual international ACM SIGIR conference on Research and development in information retrieval**

Publisher: British Computer Society

Full text available:  [pdf\(720.73 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#)

We present an integrated method for record clustering and reorganization which can be applied to any set of queries whose frequencies of request are known. The clustering algorithm works by splitting and merging current clusters and, furthermore, produces a new assignment of these clusters to pages in secondary storage. The reorganization algorithm is an on-line, incremental procedure for allocating the records to their new physical locations such that the number of pages swapped in and out of t ...


7 Spatial Database Clustering: Using a cluster manager in a spatial database system



Thomas Brinkhoff

November 2001 **Proceedings of the 9th ACM international symposium on Advances in geographic information systems**

Publisher: ACM Press

Full text available:  [pdf\(2.62 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

An important goal for a spatial database system is to minimize the I/O-cost of queries and other operations. One essential component to achieve this objective is the buffer manager. The placement of the spatial objects on disk pages is another important factor; a reasonable clustering of the objects helps to minimize the I/O-cost of queries. However, it is a difficult task to define and maintain an efficient clustering. In this paper, a cluster manager is proposed, which re-clusters spatial obje ...


8 iVIBRATE: Interactive visualization-based framework for clustering large datasets



Keke Chen, Ling Liu

April 2006 **ACM Transactions on Information Systems (TOIS)**, Volume 24 Issue 2

Publisher: ACM Press

Full text available:  [pdf\(4.48 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

With continued advances in communication network technology and sensing technology, there is astounding growth in the amount of data produced and made available through cyberspace. Efficient and high-quality clustering of large datasets continues to be one of the most important problems in large-scale data analysis. A commonly used methodology for cluster analysis on large datasets is the three-phase framework of sampling/summarization, iterative cluster analysis, and disk-labeling. There are th ...

Keywords: Clustering, interactive visualization, labeling, large datasets, performance

9 On the performance of object clustering techniques



Manolis M. Tsangaris, Jeffrey F. Naughton

June 1992 **ACM SIGMOD Record , Proceedings of the 1992 ACM SIGMOD international conference on Management of data SIGMOD '92**, Volume 21 Issue 2

Publisher: ACM Press

Full text available: pdf(1.20 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We investigate the performance of some of the best-known object clustering algorithms on four different workloads based upon the tektronix benchmark. For all four workloads, stochastic clustering gave the best performance for a variety of performance metrics. Since stochastic clustering is computationally expensive, it is interesting that for every workload there was at least one cheaper clustering algorithm that matched or almost matched stochastic clustering. Unfortunately, for each workl ...

10 Self-adaptive, on-line reclustering of complex object data



William J. McIver, Roger King

May 1994 **ACM SIGMOD Record , Proceedings of the 1994 ACM SIGMOD international conference on Management of data SIGMOD '94**, Volume 23 Issue 2

Publisher: ACM Press

Full text available: pdf(1.19 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

A likely trend in the development of future CAD, CASE and office information systems will be the use of object-oriented database systems to manage their internal data stores. The entities that these applications will retrieve, such as electronic parts and their connections or customer service records, are typically large complex objects composed of many interconnected heterogeneous objects, not thousands of tuples. These applications may exhibit widely shifting usage patterns due to their i ...

11 Data placement in Bubba



George Copeland, William Alexander, Ellen Boughter, Tom Keller

June 1988 **ACM SIGMOD Record , Proceedings of the 1988 ACM SIGMOD international conference on Management of data SIGMOD '88**, Volume 17 Issue 3

Publisher: ACM Press

Full text available: pdf(1.41 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper examines the problem of data placement in Bubba, a highly-parallel system for data-intensive applications being developed at MCC. "Highly-parallel" implies that load balancing is a critical performance issue. "Data-intensive" means data is so large that operations should be executed where the data resides. As a result, data placement becomes a critical performance issue. In general, determining the optimal placement of d ...

12 Multi-level hierarchies for scalable ad hoc routing

Elizabeth M. Belding-Royer

September 2003 **Wireless Networks**, Volume 9 Issue 5

Publisher: Kluwer Academic Publishers

Full text available: pdf(465.16 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Ad hoc networks have the notable capability of enabling spontaneous networks. These networks are self-initializing, self-configuring, and self-maintaining, even though the underlying topology is often continually changing. Because research has only begun to scratch the surface of the potential applications of this technology, it is important to prepare for the widespread use of these networks. In anticipation of their ubiquity, the protocols designed for these networks must be scalable. This inc ...

Keywords: ad hoc networks, hierarchial routing, mobile networking, scalability

13 Optimization of a hierarchical file organization for spelling correction



Tetsuro Ito, Clement T. Yu

June 1985 **Proceedings of the 8th annual international ACM SIGIR conference on Research and development in information retrieval**

Publisher: ACM Press

Full text available: pdf(511.78 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

A spelling program using a hierarchically organized file seems to be promising, since it can correct more than common typing mistakes. However, its speed of detecting spelling errors in the inputs is rather slow. Here some techniques of modifying the program to improve the speed are presented.

14 Towards an efficient management of objects in a distributed environment



A. El Habbash, J. Grimson, C. Horn

July 1990 **Proceedings of the second international symposium on Databases in parallel and distributed systems**

Publisher: ACM Press

Full text available: pdf(1.01 MB) Additional Information: [full citation](#), [references](#), [index terms](#)

15 Fast object partitioning using Stochastic learning automata



B. J. Oommen, D. Ma

November 1987 **Proceedings of the 10th annual international ACM SIGIR conference on Research and development in information retrieval**

Publisher: ACM Press

Full text available: pdf(935.94 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Let $\mathcal{O} = \{A_1, \dots, A_W\}$ be a set of W objects to be partitioned into R classes $\{P_1, \dots, P_R\}$. The objects are accessed in groups of unknown size and the size of these groups need not be equal. Additionally, the joint access probabilities of the objects are unknown. The intention is that the objects accessed more frequently together are located in the same class. This problem has been shown to be ...

16 A framework for effective retrieval



C. T. Yu, W. Meng, S. Park

June 1989 **ACM Transactions on Database Systems (TODS)**, Volume 14 Issue 2

Publisher: ACM Press

Full text available: pdf(1.56 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

The aim of an effective retrieval system is to yield high recall and precision (retrieval effectiveness). The nonbinary independence model, which takes into consideration the number of occurrences of terms in documents, is introduced. It is shown to be optimal under the assumption that terms are independent. It is verified by experiments to yield significant improvement over the binary independence model. The nonbinary model is extended to normalized vectors and is applicable to more genera ...

17 Index configuration in object-oriented databases

Elisa Bertino

July 1994 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 3 Issue 3

Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(2.23 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

In relational databases, an attribute of a relation can have only a single primitive value,

making it cumbersome to model complex objects. The object-oriented paradigm removes this difficulty by introducing the notion of nested objects, which allows the value of an object attribute to be another object or a set of other objects. This means that a class consists of a set of attributes, and the values of the attributes are objects that belong to other classes; that is, the definition of a class fo ...

Keywords: index selection, physical database design, query optimization

18 Poster papers: Learning to match and cluster large high-dimensional data sets for data integration



William W. Cohen, Jacob Richman

July 2002 **Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining**

Publisher: ACM Press

Full text available: [pdf\(634.07 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Part of the process of data integration is determining which sets of identifiers refer to the same real-world entities. In integrating databases found on the Web or obtained by using information extraction methods, it is often possible to solve this problem by exploiting similarities in the textual names used for objects in different databases. In this paper we describe techniques for clustering and matching identifier names that are both scalable and *adaptive*, in the sense that they can ...

Keywords: clustering, large datasets, learning, text mining

19 Unsupervised anomaly detection in network intrusion detection using clusters

Kingsly Leung, Christopher Leckie

January 2005 **Proceedings of the Twenty-eighth Australasian conference on Computer Science - Volume 38 ACSC '05**

Publisher: Australian Computer Society, Inc.

Full text available: [pdf\(272.72 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#)

Most current network intrusion detection systems employ signature-based methods or data mining-based methods which rely on labelled training data. This training data is typically expensive to produce. Moreover, these methods have difficulty in detecting new types of attack. Using unsupervised anomaly detection techniques, however, the system can be trained with unlabelled data and is capable of detecting previously "unseen" attacks. In this paper, we present a new density-based and grid-based cl ...

20 PageCluster: Mining conceptual link hierarchies from Web log files for adaptive Web site navigation



Jianhan Zhu, Jun Hong, John G. Hughes

May 2004 **ACM Transactions on Internet Technology (TOIT)**, Volume 4 Issue 2

Publisher: ACM Press

Full text available: [pdf\(280.84 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

User traversals on hyperlinks between Web pages can reveal semantic relationships between these pages. We use user traversals on hyperlinks as weights to measure semantic relationships between Web pages. On the basis of these weights, we propose a novel method to put Web pages on a Web site onto different conceptual levels in a link hierarchy. We develop a clustering algorithm called PageCluster, which clusters conceptually-related pages on each conceptual level of the link hierarchy based on th ...

Keywords: Link hierarchies, Web site navigation, bibliographic analysis, clustering, conceptual link hierarchies, link similarity

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads: [!\[\]\(d3fb9f94af8b26d1c844efa9a98805b0_img.jpg\) Adobe Acrobat](#) [!\[\]\(78eb1652b591ce460bbb1a853a52e223_img.jpg\) QuickTime](#) [!\[\]\(73569cd2fc0daa66f7f79a4aa32515bb_img.jpg\) Windows Media Player](#) [!\[\]\(b12ef55cf189ee2bc15774f5ac32d31a_img.jpg\) Real Player](#)